NOνA Database Requirements

Jonathan M. Paley, Julie Trumbo, Gerald Guglielmo,
Igor Mandrichenko, Andrew Norman

February 10, 2007

1 Introduction

The proposed NOνA experiment will consist of two functionally identical detectors, a "near" detector, located at FNAL and a "far" detector, located approximately 810 km away in northern Minnesota. The data acquisition (DAQ) system for each detector will require a database connection to start a data-collection run. Experience from the MINOS experiment, which also has a far detector located in nothern Minnesota, tells us that the link between the far detector and the FNAL network is tenuous, and connections are often dropped for durations of several tens of seconds. Furthermore, the far detector has over 5×10^5 read-out channels, and the expected average write-rate to the database is expected to be on the order of 10 kB/s (dominated by the DAQ and DCS systems). Therefore both far and near detectors will require a separate database server. The database server located at FNAL will however likely serve as a "central" server for online monitoring summary tables for both detectors, as well as the hardware and offline calibration and alignment tables.

During the early stages (eg, R&D, construction) of the NOνA project, measured properties of the various components of the detectors, as well as data on the progress of the construction of each detector will be stored in the NOνA database via a a web interface. Database tables that hold these data will be refered to as "hardware tables". As soon as detector modules are installed during the two year long construction period, they will become operational and will collect data. Therefore, the performance of the NOνA detectors will be tracked not only after the detector is fully constructed and is fully operational, but also during the construction period. Database tables that hold data on the operational status of the detectors will be refered to as "operational tables".

We assume that data from both hardware tables and operational tables will be used in the offline analysis of the data. Offline analysis is expected

to occur during operations, and in fact some of the data analysis will feed back in to the operations tables (eg, calibration constansts for the front-end board (FEB) electronics).

This document will outline the hardware and software requirements for the NOνA database server, which will host both the online and offline databases. The online database refers to the database that holds the hardware and operational tables. The offline database refers to the database tables that will be used in the offline reconstruction and analysis of the data. Note that the online and offline tables may or may not be held in the same database (that is, there is no requirement that all database tables be held in the same database).

2 Hardware Requirements

The requirements for the NOνA database server hardware are as follows:

1. Detector database hardware will be positioned in MN and at Fermilab.
2. Both database servers must operate with a 99% uptime.
3. 24/7 technical support for hardware is required for both near and far detector database servers.
4. Hardware failover should be implemented for both the near and far detector locations, by providing a backup system available for use.
5. The database hardware and OS must be chosen to support the expected performance and capacity requirements, must support the chosen database platform, support the operational uptime and reliability requirements of the experiment during data taking at both the MN and FNAL sites.
6. DOE and Fermilab mandated security baselines will be followed for the database hardware.

3 Software Requirements

The requirements for the NOνA database server software are as follows:

1. The database for NOνA must be a robust, reliable engine capable of operating at:
   (a) 99% uptime when database machine is up and beam is present
   (b) 95% uptime when database machine is up and beam in not present
2. The far detector database server must be able to handle a write rate of 10 kB/s.
3. The far detector database must scale to approximately 100 GB. This leaves a safe margin for the 30 GB of transient stortage for DAQ monitoring. The DCS system will not require a transient database at either detector site.
4. The near detector database must scale to 500 GB. We believe this is sufficient storage for several years' worth of data to be stored in this database.
5. Data will need to be available for 20 years. However, most of the data may be moved to long-term storage after some time period (eg, 1-2 years of running).
6. No data loss of data stored in the database due to disk failure.
7. No data loss of data stored in the database due to an operating system crash or hang.
8. No more than 1% of data loss for any reason of data already stored in the database.
9. DOE and Fermilab mandated security baselines will be followed for both the operating system and database.
10. Regularly scheduled downtimes will be allowed for patching. Minimally these will be available quarterly in sync with the security patch releases for the operating system and database.
11. Non quarterly downtimes will be granted for emergency patching as required.